Update

Advances in Processors, Coprocessors, and Accelerators in High-Performance Computing: HPC User Forum, April 7–9, 2014, Santa Fe, New Mexico

Earl C. Joseph, Ph.D.
Steve Conway
Chirag Dekate, Ph.D.

IN THIS UPDATE

This IDC update captures part of the proceedings at the 52nd HPC User Forum held in Santa Fe, New Mexico. The topic covered in detail in this document is innovations in processor, accelerator, and coprocessor technologies targeting the high-performance computing (HPC) ecosystem. This update is a near-verbatim record of the talks given in this session.

"Micron's Automata Processor": Paul Dlugosch

We were founded in 1978. Our FY13 revenue was $9 billion, and we're running twice that fast in FY14. We have 30,000 employees worldwide. All of our products are memory storage devices: DRAM, NAND flash, SSDs, bare dies, and more.

The five big tech trends driving our business are networking, machine to machine, mobile, cloud, and big data. Big data needs high-memory performance. Power is as important in HPC as in the mobile market. For 30 years, our customers have said memory is the problem and they need faster memory. The memory industry has responded, "Sure, we can do that."

Our Hybrid Memory Cube was launched last year. We believe that the demands of HPC and big data will be well served by the Automata processor (AP). Micron is not actively promoting the Automata processor yet. It is a massive array of very fine-grained processing elements. These are very simple elements with a very complicated interconnect. Version 1.4 of the SDK is available now. We are evaluating early markets.

Application examples include network security for QoS and deep packet inspection. Sample availability of the processor is targeted for 3Q14. It has a 2D fabric of about 50,000 processing elements built on a 20nm die.

CPUs have been asked to do every type of problem and don't do all of them well. Network security, bioinformatics, video analytics, and data analytics are targeted apps for the processor. They represent themselves in the form of graphs. Today, we transform these problems into a form that will run on CPUs, and that causes problems for programmer productivity.

The Automata processor architecture is designed to handle these kinds of problems. Its architecture is akin to that of an artificial neural network. The processor directly takes data streams and drives them into a row decoder. For programming, we developed the AP Workbench for hierarchical design and macro support. Parallelization requires no special consideration by the user. Each automaton operates independently on the input data stream.

"The IBM-DOME 64-Bit Microserver Demonstrator: Findings, Status, and Outlook": Ronald Luijten, IBM Zurich

Compute is free – data is not. The data revolution is only starting. My definition of a microserver is integrating an entire server node motherboard into a single microchip, except for DRAM, NOR boot flash, and power conversion logic.

I got funding from the SKA project, whose goal is to measure the Big Bang. This will be placed in the Australian desert and another location. In the desert, there is no power grid and no Internet, so you need dirt-cheap computing. It would need 80-1,200 exaflops.

The IBM-ASTRON DOME project includes seven projects for technology road map development for the SKA. The goal is to create the world's highest-density 64-bit microserver drawer for SKA and future IBM business: a datacenter in a box, with very high-energy efficiency and very low cost, using commodity components.

Another challenge is to leverage IBM cooling technology (e.g., SuperMUC). It must also be true 64-bit to enable business applications. This is a research project. The strongest ecosystem today is x86 and the fastest growing is ARM.

Status: The Rev 2 board is in bringup, the Rev 1 board is being populated, the power module is being debugged, the multinode carrier board is in bringup, and the water cooling is coming along well.

"ARM Processor Directions": Dwight Barron, HP

I'm chief technologist for the HP Moonshot group. Our HP Lab research related to this dates back around 10 years. For the past 5 years, I've worked in the hyperscale business that includes HPC and Internet providers. We concluded that energy-efficient computing would be driven more by throughput than by single-threaded performance. For this, we looked to mobile devices, where each new generation has gotten more work per watt. We morphed this processor into a server.

To get energy efficiency, the client world uses a purpose-built device. You ask what it does, not what it is. In mobile devices, die real estate is split about 50:50 between graphics and integer, with the graphics part growing.

When we translated this to servers on the Moonshot side, we knew we needed to ask what functions the server needs to perform and how to integrate that into the silicon. Various types of companies deeply understand various functions, such as DSP companies understanding signal processing. The microprocessor progression today is a new architecture every 3 years and a refresh every 18 months.

The first goal for Moonshot was that it had to be a purpose-built, functional mentality. We will talk about what each server does, such as "serve a million Web pages." So, in Moonshot, we are moving from very large general-purpose systems to purpose-built systems, and we will ask customers what they want to be able to do, that is, what their problems are.

"A Rare Look at Real-World Data Analysis of Supercomputer Faults — DRAM, SRAM, and GPGPUs": Nathan DeBardeleben, LANL

This is a collaboration with vendors and universities. The systems include Cray Cielo, Jaguar, and Hopper supercomputers as well as the Moonlight GPU cluster at LANL. We look at corrected and uncorrected faults in DRAM and SRAM, altitude effects, positional effects, and GPGPU variability in memory bandwidth, performance, and reliability.

Cielo affords us lots of data, and not much of it has been published. Our report is the largest ever done – on 24 billion DRAM hours. Our report identifies problems in others' analyses and how these have led to wrong conclusions. For example, we showed that DRAM reliability varies by vendor. There were differences of as much as four times in reliability.

Study findings include:

- Fault rates increase 20% as you rise vertically in a rack, as we saw on both Jaguar and Cielo. We haven't identified the reason for these correctable errors.
- DDR command and address parity is a good thing.
- SRAM memory is affected by altitude in the form of increased fault rates due to increased cosmic ray-induced neutron strikes. There may be additional effects, such as alpha particles. Conclusion: Most SRAM faults in the field are caused by known high-energy particles.
- The best way to reduce SRC uncorrected error rates is to extend single-bit correction (e.g., ECC) through additional structures in the processor.

We have done predictions for exascale SRAM and DRAM. We expect a 5–90 times increase in uncorrected errors – 5 times would be tolerable, while 90 times would not.

We performed the GPGPU field study using the Moonlight GPGPU cluster at LANL. It's a Fermi-generation system deployed in 2011 and has 308 nodes and dual socket, eight-core Sandy Bridge processors, with two NVIDIA Tesla M2090 GPGPUs per node. This was LANL's first production foray into GPGPU computing. The problems were mainly in the evolving software and driver stack. These problems were fixed and are not representative of current-generation GPGPUs.

"Accelerating Physics-Based Seismic Hazard Analysis on Heterogeneous Supercomputers": Yifeng Cui

Dr. Yifeng Cui, High Performance Geocomputing Laboratory at the University of California, San Diego, was a winner at SC13. This achievement was by a large consortium of more than 60 United States-based and international institutions.

This project is aimed at helping advance earthquake prediction using a GPU-enabled code. Computational requirements for earthquake simulations have increased exponentially. Vendor supporters are IBM, Cray, and NVIDIA. We developed the AWD-ODC code for use with a hybrid CPU-GPU architecture. The most important model in our achievement is the probabilistic seismic hazard model. It includes an earthquake rupture forecast, attenuation relationship, and intensity measures.

The ultimate goal is to create a broadband CyberShake hazard model for all of California (1,400 sites). Some of this work has been done, and some remains to be done by a study proposed for 2015. The CyberShake SGT work is being done on a Cray XK7. We have looked ahead as far as exascale and what that will make possible for this work. The goal is to be able to complete the challenging simulation within 24 hours. In the next two years, we plan to apply 100PF to this problem.

LEARN MORE

Related Research

Additional research from IDC in the technical computing hardware program includes the following documents:

- *Worldwide High-Performance Data Analysis 2013–2017 Forecast* (IDC #241315, June 2013)
- *Experiences in HPC: HPC User Forum, April 29–May 1, 2013, Tucson, Arizona* (IDC #241455, June 2013)
- *HPC in Aerospace, Astrophysics, and Astronomy: HPC User Forum, April 29–May 1, 2013, Tucson, Arizona* (IDC #241451, June 2013)
- *Top Issues for HPC Sites: HPC User Forum, April 29–May 1, 2013, Tucson, Arizona* (IDC #241463, June 2013)
- *Potential Disruptive Technologies Panel: HPC User Forum, April 29–May 1, 2013, Tucson, Arizona* (IDC #241452, June 2013)
- *Advanced Visualization: HPC User Forum, April 29–May 1, 2013, Tucson, Arizona* (IDC #241446, June 2013)
- *Worldwide Technical Computing Server 2013–2017 Forecast* (IDC #241154, May 2013)
- *Supercomputers Exceed 50% of the HPC Server Market in 2012* (IDC #240426, April 2013)
- *Changing Market Dynamics: HPC Meeting Big Data and IDC's Projected Evolution of the Market* (IDC #240365, March 2013)
- *Livermore Lab Expands Industry Partnerships: Economic Security Is Vital for National Security* (IDC #240232, March 2013)
- *Worldwide Technical Computing 2013 Top 10 Predictions* (IDC #239421, February 2013)
- *High-Performance Data Analysis at NASA JPL* (IDC #238254, December 2012)
- *Advanced Research at the South Ural State University Supercomputing Center* (IDC #238225, December 2012)

- *The Economic Value of HPC in Science and Industry: HPC User Forum, September 2012, Dearborn, Michigan* (IDC #237182, October 2012)
- *Big Data in HPC: HPC User Forum, September 2012, Dearborn, Michigan* (IDC #237180, October 2012)
- *How Nations Are Applying High-End Petascale Supercomputers for Innovation and Economic Advancement in 2012* (IDC #236341, August 2012)
- *Tokyo Institute of Technology: Global Scientific Information and Computing Center* (IDC #236243, August 2012)
- *Shanghai Supercomputer Center* (IDC #236245, August 2012)
- *Petascale Supercomputing at the University of Tokyo* (IDC #236292, August 2012)
- *HPC Application Leadership at the Supercomputing Center of Chinese Academy of Sciences* (IDC #236281, August 2012)
- *India Broadening Access to Supercomputing* (IDC #IcUS23621912, July 2012)
- *Institute of Process Engineering, Chinese Academy of Sciences* (IDC #236204, July 2012)
- *How Fujitsu Built the World's Fastest Supercomputer in Record Time and Ahead of Schedule* (IDC #235733, July 2012)
- *Europe Sharpens Its Focus on Exascale Computing* (IDC #IcUS23555112, June 2012)
- *IBM Returns Supercomputer Crown to United States* (IDC #IcUS23547812, June 2012)
- *Potential Disruptive Technologies: HPC User Forum, April 2012, Richmond, Virginia* (IDC #234742, May 2012)
- *The Broader HPC Market: Servers, Storage, Software, Middleware, and Services* (IDC #234682, May 2012)
- *HPC End-User Site Update: RIKEN Advanced Institute for Computational Science* (IDC #233690, March 2012)
- *National Supercomputing Center in Tianjin* (IDC #233971, March 2012)

- *HPC End-User Site Update: RIKEN Advanced Institute for Computational Science* (IDC #233690, March 2012)
- *National Supercomputing Center in Tianjin* (IDC #233971, March 2012)

- *Worldwide Data Intensive-Focused HPC Server Systems 2011-2015 Forecast* (IDC #232572, February 2012)

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-insights-community.com
www.idc.com